

## CLAIMS

- 1 1. A method for a coordinated bringup of a repaired storage appliance in a storage  
2 appliance cluster, the repaired storage appliance having a disk subsystem, the method  
3 comprising the steps of:
  - 4 asserting a GIVEWAIT state in a predetermined memory location of the repaired  
5 storage appliance;
  - 6 releasing disk reservations in response to detection of the asserted GIVEWAIT  
7 state by a surviving storage appliance;
  - 8 initializing the disk subsystem of the repaired storage appliance;
  - 9 asserting a MBWAIT state in the predetermined memory location; and
  - 10 performing a giveback operation by the surviving storage appliance in response to  
11 detecting the MBWAIT state.
- 1 2. The method of claim 1 further comprising the steps of:
  - 2 completing the repaired storage appliance initialization; and
  - 3 processing data access requests by the repaired storage appliance.
- 1 3. The method of claim 1 wherein the predetermined memory location comprises a  
2 state data structure within a memory of the repaired storage appliance.
- 1 4. The method of claim 1 wherein the surviving storage appliance detects the  
2 GIVEWAIT state by performing a remote direct memory access read operation to the  
3 predetermined memory location.
- 1 5. The method of claim 1 wherein the surviving storage appliance detects the  
2 MBWAIT state by performing a remote direct memory access operation of the predeter-  
3 mined memory location.

1 6. The method of claim 1 wherein the surviving storage appliance ceases to process  
2 data access requests directed to the repaired storage appliance after performing the give-  
3 back operation.

1 7. A storage appliance for use in a storage system cluster, the storage appliance  
2 comprising:

3 a storage operating system having a cluster failover layer adapted to perform a  
4 coordinated bringup operation in association with a partner storage appliance, wherein  
5 the coordinated bringup operation comprises the steps of:

- 6 (i) asserting a first state in a predetermined memory location of the storage  
7 appliance;
- 8 (ii) initializing a disk subsystem of the repaired storage appliance in re-  
9 sponse to detecting a release of disk reservations by a partner storage appliance;
- 10 (iii) asserting a second state in the predetermined memory location;
- 11 (iv) processing data access requests directed to the storage appliance after  
12 a giveback operation performed by the partner storage appliance; and  
13 whereby a period of time during which clients of the storage system are without  
14 connectivity is minimized.

1 8. The storage appliance of claim 6 wherein the cluster failover layer is further  
2 adapted to perform routine remote direct and memory access read operations to the part-  
3 ner storage appliance to detect a state of the partner storage appliance.

1 9. The storage appliance of claim 8 wherein the second state comprises a MBWAIT  
2 state.

1 10. The storage appliance of claim 8 wherein the first state comprises a GIVEWAIT  
2 state.

- 1 11. A method for a coordinated bringup of a repaired storage appliance in a storage
- 2 appliance cluster, the repaired storage appliance having a disk subsystem, the method
- 3 comprising the steps of:
  - 4 asserting a first state in a predetermined memory location of the repaired storage
  - 5 appliance;
  - 6 releasing disk reservations in response to detection of the asserted first state by a
  - 7 surviving storage appliance;
  - 8 initializing the disk subsystem of the repaired storage appliance;
  - 9 asserting a second state in the predetermined memory location; and
  - 10 performing a giveback operation by the surviving storage appliance in response to
  - 11 detecting the second state.
- 12
- 1 12. The method of claim 11 wherein the predetermined memory location comprises a
- 2 state data structure within a memory of the repaired storage appliance.
- 1 13. The method of claim 11 wherein the surviving storage appliance detects the first
- 2 state by performing a remote direct memory access read operation to the predetermined
- 3 memory location.
- 1 14. The method of claim 11 wherein the surviving storage appliance detects the sec-
- 2 ond state by performing a remote direct memory access operation of the predetermined
- 3 memory location.
- 1 15. The method of claim 11 wherein the surviving storage appliance ceases to process
- 2 data access requests directed to the repaired storage appliance after performing the give-
- 3 back operation.
- 1 16. The method of claim 11 wherein the first state comprises a GIVEWAIT state.

- 1 17. The method of claim 11 wherein the second state comprises a MBWAIT state.
- 1 18. The method of claim 11 wherein the set of disk reservations comprises small  
2 computer systems interface reservations.
- 1 19. A computer readable medium, including program instructions executing on a  
2 storage appliance, for a coordinated bringup of a repaired storage appliance in a storage  
3 appliance cluster, the repaired storage appliance having a disk subsystem, the computer  
4 readable medium including instructions for performing the steps of:
  - 5 asserting a GIVEWAIT state in a predetermined memory location of the repaired  
6 storage appliance;
  - 7 releasing disk reservations in response to detection of the asserted GIVEWAIT  
8 state by a surviving storage appliance;
  - 9 initializing the disk subsystem of the repaired storage appliance;
  - 10 asserting a MBWAIT state in the predetermined memory location; and
  - 11 performing a giveback operation by the surviving storage appliance in response to  
12 detecting the MBWAIT state.
- 1 20. The computer readable medium of claim 19 further comprising the steps of:
  - 2 completing the repaired storage appliance initialization; and
  - 3 processing data access requests by the repaired storage appliance.
- 1 21. The computer readable medium of claim 19 wherein the predetermined memory  
2 location comprises a state data structure within a memory of the repaired storage appli-  
3 ance.
- 1 22. The computer readable medium of claim 19 wherein the surviving storage appli-  
2 ance detects the GIVEWAIT state by performing a remote direct memory access read  
3 operation to the predetermined memory location.

- 1 23. The computer readable medium of claim 19 wherein the surviving storage appli-
- 2 ance detects the MBWAIT state by performing a remote direct memory access operation
- 3 of the predetermined memory location.

1